

should bring an appropriate response; that response should depend on the results and should not necessarily be to adjust the parameter estimates. When realized default, severity, or exposures rates diverge from expected ranges, those divergences may point to issues in the estimation or mapping elements of quantification. They may also indicate potential problems in other parts of the ratings assignment process. The bank's validation policy must describe (at least in broad terms) the types of responses that should be considered when relevant action thresholds are crossed.

Appendix to Part III: Illustrations of the Quantification Process

This appendix provides examples to show how the logical framework described in this guidance, with its four stages (data, estimation, mapping, and application), applies when analyzing typical current bank practices. The framework is broadly applicable—for PD or LGD or EAD; using internal, external, or pooled reference data; for simple or complex estimation methods—although the issues and concerns that arise at each stage depend on a bank's approach. These examples are intended only to illustrate the logic of the four-stage IRB quantification framework, and should not be taken to endorse the particular techniques presented in the examples. In fact, certain aspects of the examples are not consistent with the standards outlined in this guidance.

Example 1: PD Estimation From Bond Data

- A bank establishes a correspondence between its internal grades and external rating agency grades; the bank has determined that its Grade 4 is equivalent to $\frac{3}{4}$ BB and $\frac{1}{4}$ B on the Standard and Poor's scale.
- The bank regularly obtains published estimates of mean default frequencies for publicly rated BB and B obligors in North America from 1970 through 2002.
- The BB and B historical default frequencies are weighted 75/25, and the result is a preliminary PD for the bank's internal Grade 4 credits.
- However, the bank then increases the PD by 10 percent to account for the fact that the S&P definition of default is more lenient than the IRB definition.
- The bank makes a further adjustment to ensure that the resulting grade PD is greater than the PD attributed to Grade 3 and less than the PD attributed to Grade 5.
- The result is the final PD estimate for Grade 4.

Process Analysis for Example 1

Data—The reference data set consists of issuers of publicly rated debt in North America over the period 1970 through 2002. The data description is very basic: each issuer in the reference data is described only by its rating (such as AAA, AA, A, BBB, and so on).

Estimation—The bank could have estimated default rates itself using a database purchased from Standard and Poor's, but

since these estimates would just be the mean default rates per year for each grade, the bank could just as well (and in this example does) use the published historical default rates from S&P; in essence, the estimation step has been outsourced to S&P. The 10 percent adjustment of PD is part of the estimation process in this case because the adjustment was made prior to the application of the agency default rates to the internal portfolio data.

Mapping—The bank's mapping is an example of a grade mapping; internal Grade 4 is linked to the 75/25 mix of BB and B. Based on the limited information presented in the example, this step should be explored further. Specifically, how did the bank determine the 75/25 mix?

Application—Although the application step is relatively straightforward in this case, the bank does make the adjustment of the Grade 4 PD estimate to give it the desired relationship to the adjacent grades. This adjustment is part of the application stage because it is made after the adjusted agency default rates are applied to the internal grades.

Example 2: PD Estimation Using a Merton-Type Equity-Based Model

- A bank obtains a 20-year database of North American firms with publicly traded equity, some of which defaulted during the 20-year period.
- The bank uses the Merton approach to modeling equity in these firms as a contingent claim, constructing an estimate of each firm's distance-to-default at the start of each year in the database. The bank then ranks the firm-years within the database by distance-to-default, divides the ordered observations into 20 equal groups or buckets, and computes a mean historical one-year default frequency for each bucket. That default frequency is taken as an estimate of the applicable PD for any obligor within the range of distance-to-default values represented by each of the 20 buckets.
- The bank next looks at all obligors with publicly traded shares within each of its internal grades, applies the same Merton-type model to compute distance-to-default at quarter-end, sorts these observations into the 20 buckets from the previous step, and assigns the corresponding PD estimate.
- For each internal grade, the bank computes the mean of the individual obligor default probabilities and uses that average as the grade PD.

Process Analysis for Example 2

Data—The reference data set consists of the North American firms with publicly traded equity in the acquired database. The reference data are described in this case by a single variable, specifically an identifier of the specific distance-to-default range from the Merton model (one of the 20 possible in this case) into which a firm falls in any year.

Estimation—The estimation step is simple: the average default rate is calculated for each distance-to-default bucket. Since the data cover 20 years and a wide range of economic conditions, the resulting estimates satisfy the long-run average requirement.

Mapping—The bank maps selected portfolio obligors to the reference data set

using the distance-to-default generated by the Merton model. However, not all obligors can be mapped, since not all have traded equity. This introduces an element of uncertainty into the mapping that requires additional analysis by the bank: were the mapped obligors representative of other obligors in the same grade? The bank would need to demonstrate comparability between the publicly traded portfolio obligors and those not publicly traded. It may be appropriate for the bank to make conservative adjustments to its ultimate PD estimates to compensate for the uncertainty in the mapping. The bank also would need further analysis to demonstrate that the implied distance-to-default for each internal grade represented long-run expectations for obligors assigned to that grade; this could involve computing the Merton model for portfolio obligors over several years of relevant history that span a wide range of credit conditions.

Application—The final step is aggregation of individual obligors to the grade level through calculation of the mean for each grade, and application of this grade PD to all obligors in the grade. The bank might also choose to modify PD assignments further at this stage, combining PD estimates derived from other sources, applying adjustments for cyclicalities, introducing an appropriate degree of conservatism, or making other adjustments.

Example 3: LGD Estimation From Internal Default Data

- For each loan in its portfolio, a bank records collateral coverage as a percentage, as well as which of four types of collateral applies.
- A bank has retained data on all defaulted loans since 1995. For each defaulted loan in the database, the bank has a record of the collateral type within the same four broad categories. However, collateral coverage is only recorded at three levels (low, moderate, or high, depending on the ratio of collateral to exposure at default).
- The bank also records the timing and discounted value of recoveries net of workout costs for each defaulted loan in the database. Cash flows are tracked from the date of default to a "resolution date," defined as the point at which the remaining balance is less than 5 percent of the exposure at the time of default. A recovery percentage is computed, equal to the value of recoveries discounted to the date of default, divided by the exposure at default.
- For each cell (each of the 12 combinations of collateral type and coverage), the bank computes a simple mean LGD percentage as the mean of one minus the recovery percentage. One of the categories has a mean LGD of less than zero (recoveries have exceeded exposure on average), so the bank sets the LGD at zero to be conservative.
- The bank assigns an estimate of expected LGD to each loan in the current portfolio by using collateral information to slot it into one of the 12 cells. The bank then applies the mean historical LGD for that cell and adjusts the result upward by 10 percent to compensate for the fact that the loss data come from a period believed to be unusually good economic performance.